IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

ATTORNEY DOCKET NO. **AUS919970761US2**

In re Application of:	§	
	§	
JAMES W. ARENDT, ET AL.	§	
	§	
Serial No. **UNKNOWN**	§	Examiner: **UNKNOWN**
	§	
Filed: **JULY 10, 1998**	§	
	§	Art Unit: **UNKNOWN**
For: **HIGHLY SCALABLE AND HIGHLY AVAILABLE CLUSTER SYSTEM MANAGEMENT SCHEME**	§	

4065P4E 09/896,048

PRELIMINARY AMENDMENT

Assistant Commissioner for Patents
Washington, D.C. 20231

Sir:

Please amend the above-identified application as follows:

CERTIFICATE OF MAILING BY "EXPRESS MAIL" UNDER 37 CFR § 1.10

"Express Mail" mailing label number: ET657447168US

Date of Mailing: January 7, 2002

I hereby certify that the application/correspondence attached hereto is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to: Assistant Commissioner for Patents, Box Patent Application, Washington, D.C. 20231.

Dena Faris
Typed name of person mailing paper or fee

Dena Faris
Signature of person mailing paper or fee

PENDING CLAIMS

Claims 1-13, 21-22 and 26 have been cancelled, Claims 14 and 23 have been amended and new Claims 27-31 have been added as follows:

1. (Deleted)

2. (Deleted)

3. (Deleted)

4. (Deleted)

5. (Deleted)

6. (Deleted)

7. (Deleted)

8. (Deleted)

9. (Deleted)

10. (Deleted)

11. (Deleted)

12. (Deleted)

13. (Deleted)

14. (Amended) A method of partially replicating configuration information in a distributed database, comprising:

defining a subset of data processing systems within a cluster system as a resource group, wherein a data processing system may be a member of more than one resource group;

defining configuration data for the resource group; and

replicating the configuration data only on each data processing system within the resource group.

15. (Unchanged) The method of claim 14, wherein the step of defining a subset of data processing systems within a cluster as a resource group further comprises:

defining a highly available application and each data processing system designated to manage the application as a resource group.

16. (Unchanged) The method of claim 15, wherein the step of defining a highly available application and each data processing system managing the application as a resource group further comprises:

defining a plurality of resource groups for each highly available application within the cluster, each resource group including all data processing systems managing the corresponding application.

17. (Unchanged) The method of claim 14, wherein the step of defining configuration data for the resource group further comprises:

instantiating a configuration object containing configuration and status information for a highly available application corresponding to the resource group and having an associated list of data processing systems within the resource group.

18. (Unchanged) The method of claim 17, wherein the step of replicating the configuration data only on each data processing system within the resource group further comprises:

replicating the configuration object on each data processing system identified in an owners list associated with the configuration object.

19. (Unchanged) The method of claim 17, wherein the step of replicating the configuration data only on each data processing system within the resource group further comprises:

replicating, on a data processing system, a configuration object for each resource group including the data processing system.

20. (Unchanged) The method of claim 14, further comprising:

maintaining, on a data processing system, a configuration object for each resource group including the data processing system and no configuration objects for other resource groups.

21. (Deleted)

22. (Deleted)

23. (Amended) A computer program product in a computer usable medium, comprising:

instructions defining a subset of data processing systems within a network as a resource group, wherein a data processing system may be a member of more than one resource group;

instructions defining configuration data for the resource group; and

instructions for replicating the configuration data only on each data processing system within the resource group.

24. (Unchanged) The computer program product of claim 23, wherein the instructions defining a highly available application and each data processing system managing the application as a resource group further comprise:

instructions defining a plurality of resource groups for each highly available application within the network, each resource group including all data processing systems managing the corresponding application.

25. (Unchanged) The computer program product of claim 22, wherein the instructions defining configuration data for the resource group further comprises:

instructions instantiating a configuration object containing configuration and status information for a highly available application corresponding to the resource group and having an associated list of data processing systems within the resource group.

26. (Deleted)

27. (Newly Added) A cluster multiprocessing system, comprising:

a plurality of data processing systems segregated into a plurality of resource groups, wherein each of the plurality of data processing systems may be a member of more than one resource group;

a plurality of configuration objects each corresponding to a resource group within the plurality of resource groups; and

wherein each of the plurality of configuration objects is replicated only on each data processing system within the resource group associated with the configuration object.

28. (Newly Added) The cluster multiprocessor system of Claim 27, wherein a highly available application and each data processing system designated to manage the application is defined as a resource group.

29. (Newly Added) The cluster multiprocessor system of Claim 28, wherein a plurality of resource groups is defined for each highly available application within the cluster, each resource group including all data processing systems managing the corresponding application.

30. (Newly Added) The cluster multiprocessor system of Claim 27, wherein each of the plurality of configuration objects contains configuration and status information for a highly available application corresponding to the resource group and an associated owners list of data processing systems within the resource group.

31. (Newly Added) The cluster multiprocessor system of Claim 30, wherein each of the plurality of configuration objects are replicated on each data processing system identified in an owners list associated with the configuration object.

REMARKS

Original Claims 1-13, 21-22 and 26 have been cancelled, Claims 14 and 23 have been amended and new Claims 27-31 have been added. The newly added Claims do not introduce any new matter. Attached hereto, captioned "MARKED-UP APPENDIX," is a marked-up version of the changes made to the claims by the current amendment. Applicants respectfully request a prompt consideration of the Application. Please charge the fee of $740.00 to IBM Corporation Deposit Account No. 09-0447. No additional fee is believed to be require; however, in the event that any additional fees are required, please charge IBM Corporation Deposit Account No. 09-0447.

Respectfully submitted,

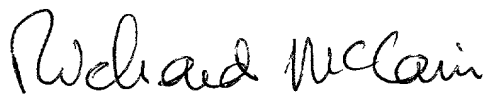

Richard N. McCain
Reg. No. 43,785
BRACEWELL & PATTERSON, LLP
Suite 350, Lakewood on the Park
7600B North Capital of Texas Highway
Austin, Texas 78731
(512) 343-6116

ATTORNEY FOR APPLICANTS

MARKED-UP APPENDIX

In the Claims:

Please delete Claims 1-13, 21-22 and 26, amend Claims 14 and 23 and add new Claims 27-31 as follows:

1. (Deleted)

2. (Deleted)

3. (Deleted)

4. (Deleted)

5. (Deleted)

6. (Deleted)

7. (Deleted)

8. (Deleted)

9. (Deleted)

10. (Deleted)

11. (Deleted)

12. (Deleted)

Copy provided by USPTO

13. (Deleted)

14. (Amended) A method of partially replicating configuration information in a distributed database, comprising:

defining a subset of data processing systems within a cluster system as a resource group, wherein a data processing system may be a member of more than one resource group;

defining configuration data for the resource group; and

replicating the configuration data only on each data processing system within the resource group.

21. (Deleted)

22. (Deleted)

23. (Amended) A computer program product in a computer usable medium, comprising:

instructions defining a subset of data processing systems within a network as a resource group, wherein a data processing system may be a member of more than one resource group;

instructions defining configuration data for the resource group; and

instructions for replicating the configuration data only on each data processing system within the resource group.

26. (Deleted)

27. (Newly Added) A cluster multiprocessing system, comprising:

a plurality of data processing systems segregated into a plurality of resource groups, wherein each of the plurality of data processing systems may be a member of more than one resource group;

a plurality of configuration objects each corresponding to a resource group within the plurality of resource groups; and

wherein each of the plurality of configuration objects is replicated only on each data processing system within the resource group associated with the configuration object.

28. (Newly Added) The cluster multiprocessor system of Claim 27, wherein a highly available application and each data processing system designated to manage the application is defined as a resource group.

29. (Newly Added) The cluster multiprocessor system of Claim 28, wherein a plurality of resource groups is defined for each highly available application within the cluster, each resource group including all data processing systems managing the corresponding application.

30. (Newly Added) The cluster multiprocessor system of Claim 27, wherein each of the plurality of configuration objects contains configuration and status information for a highly available application corresponding to an associated resource group and an associated owners list of data processing systems within the associated resource group.

31. (Newly Added) The cluster multiprocessor system of Claim 30, wherein each of the plurality of configuration objects are replicated on each data processing system identified in an owners list associated with the configuration object.

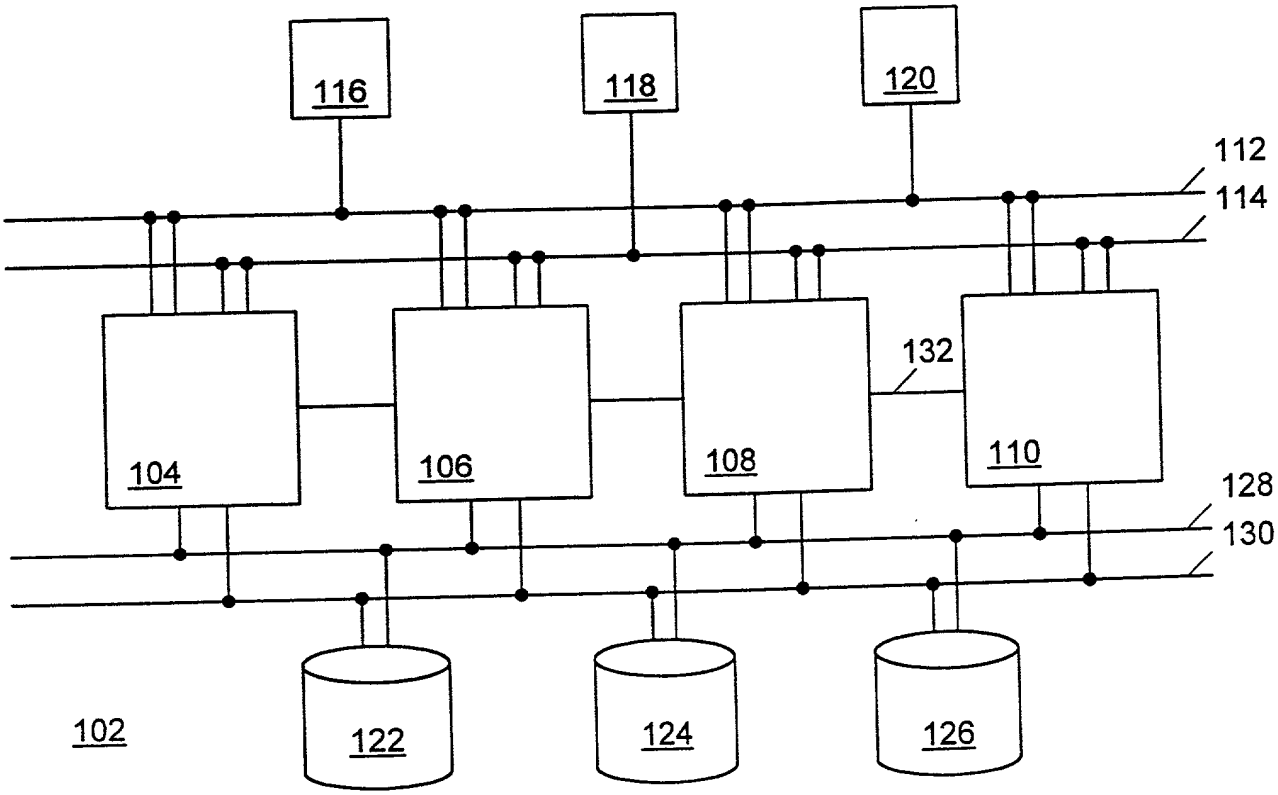

Figure 1

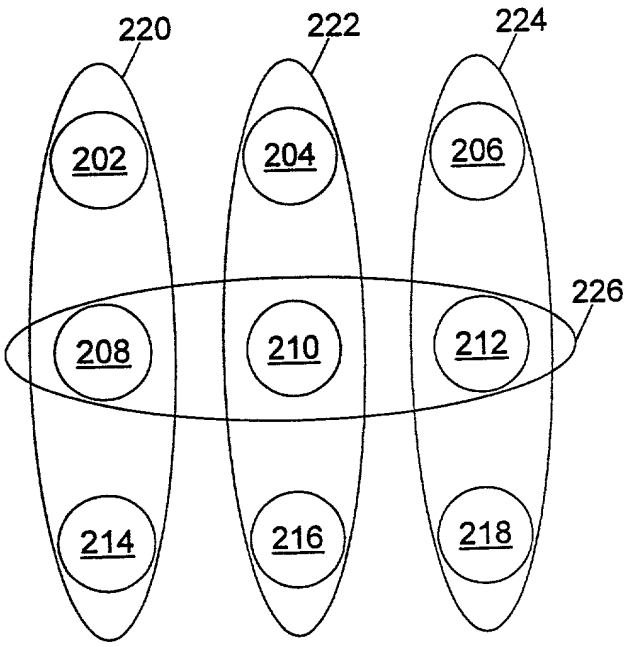

Figure 2A

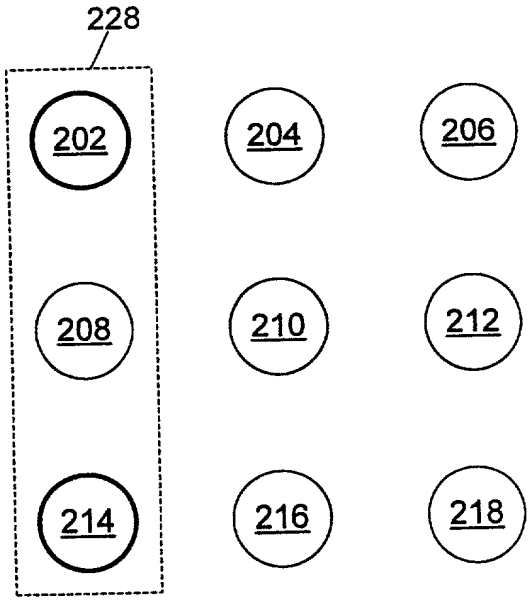

Figure 2B

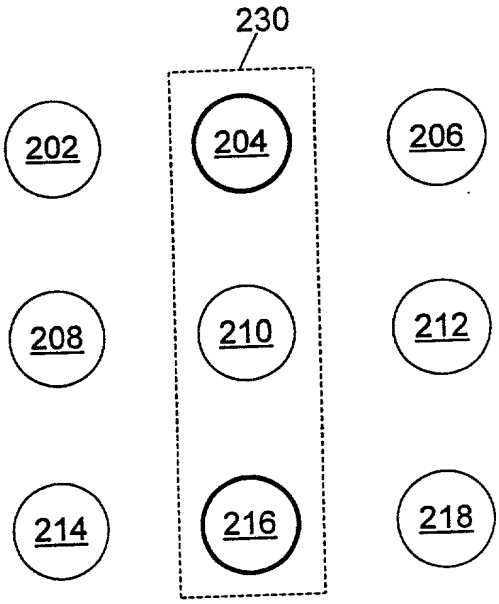

Figure 2C

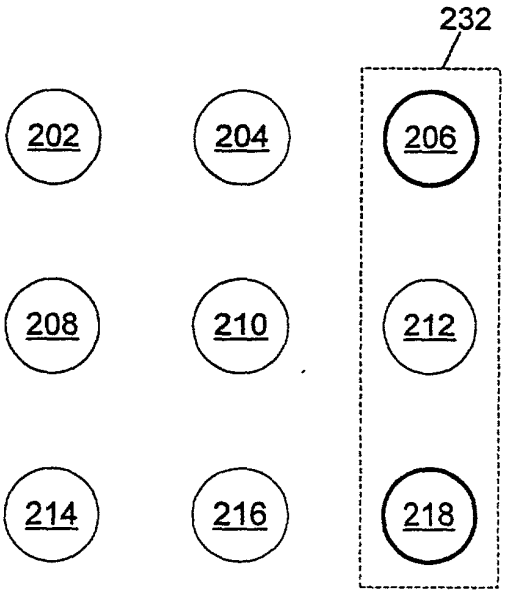

Figure 2D

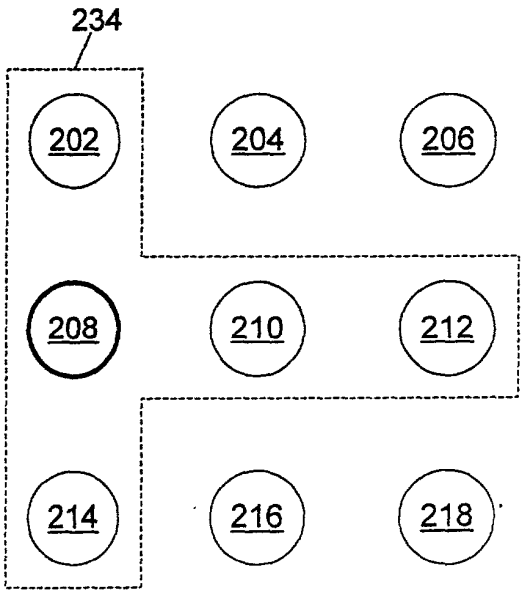

Figure 2E

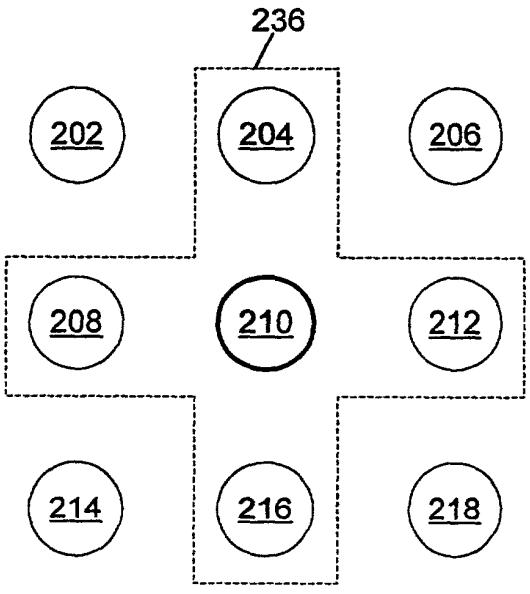

Figure 2F

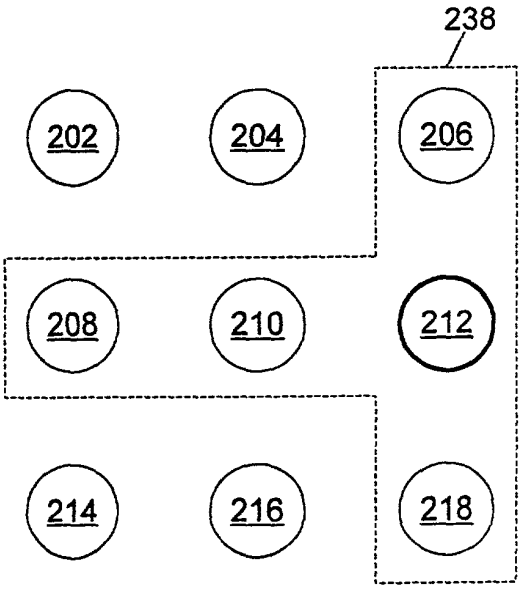

Figure 2G

EP0874302A2

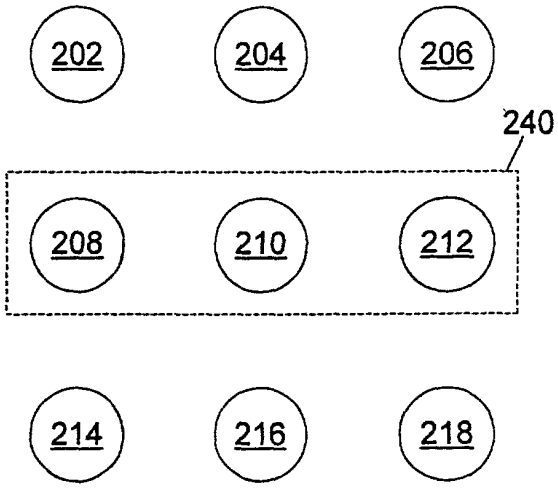

Figure 2H

EDS-2002-0293

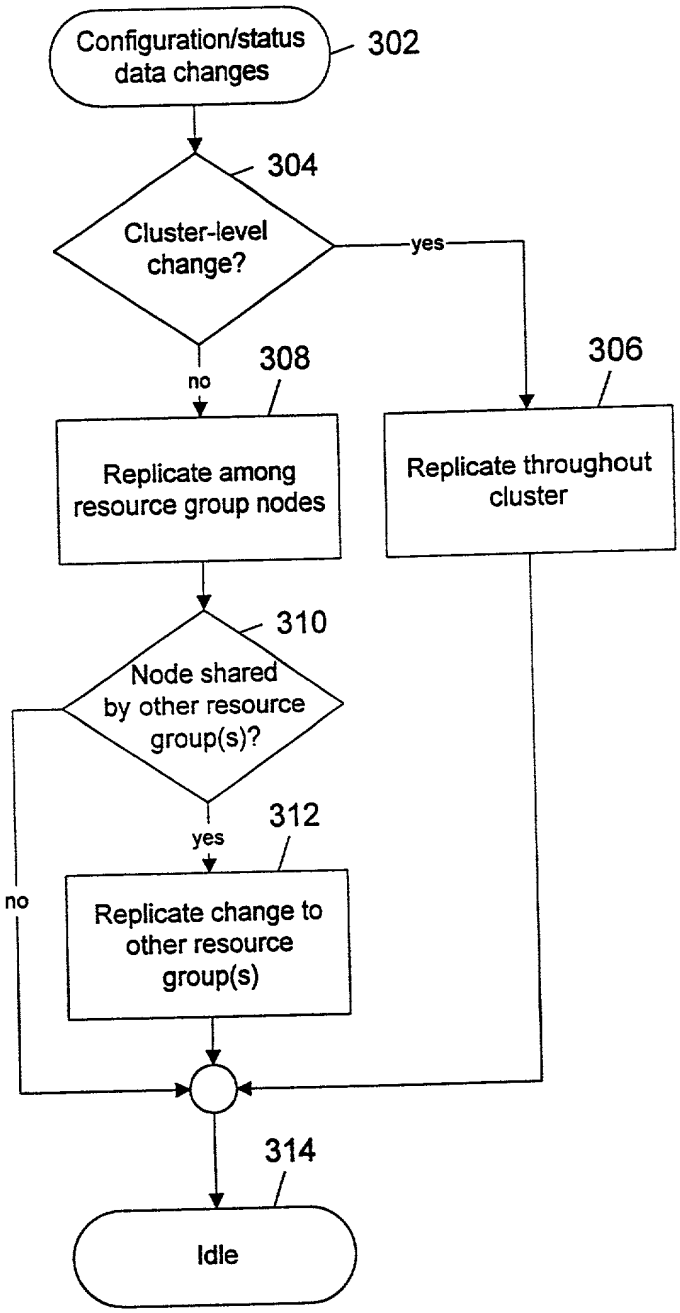
Configuration/status data changes
302
304
Cluster-level change?
yes
no
306
308
Replicate among resource group nodes
Replicate throughout cluster
310
Node shared by other resource group(s)?
yes
312
no
Replicate change to other resource group(s)
314
Idle

Figure 3

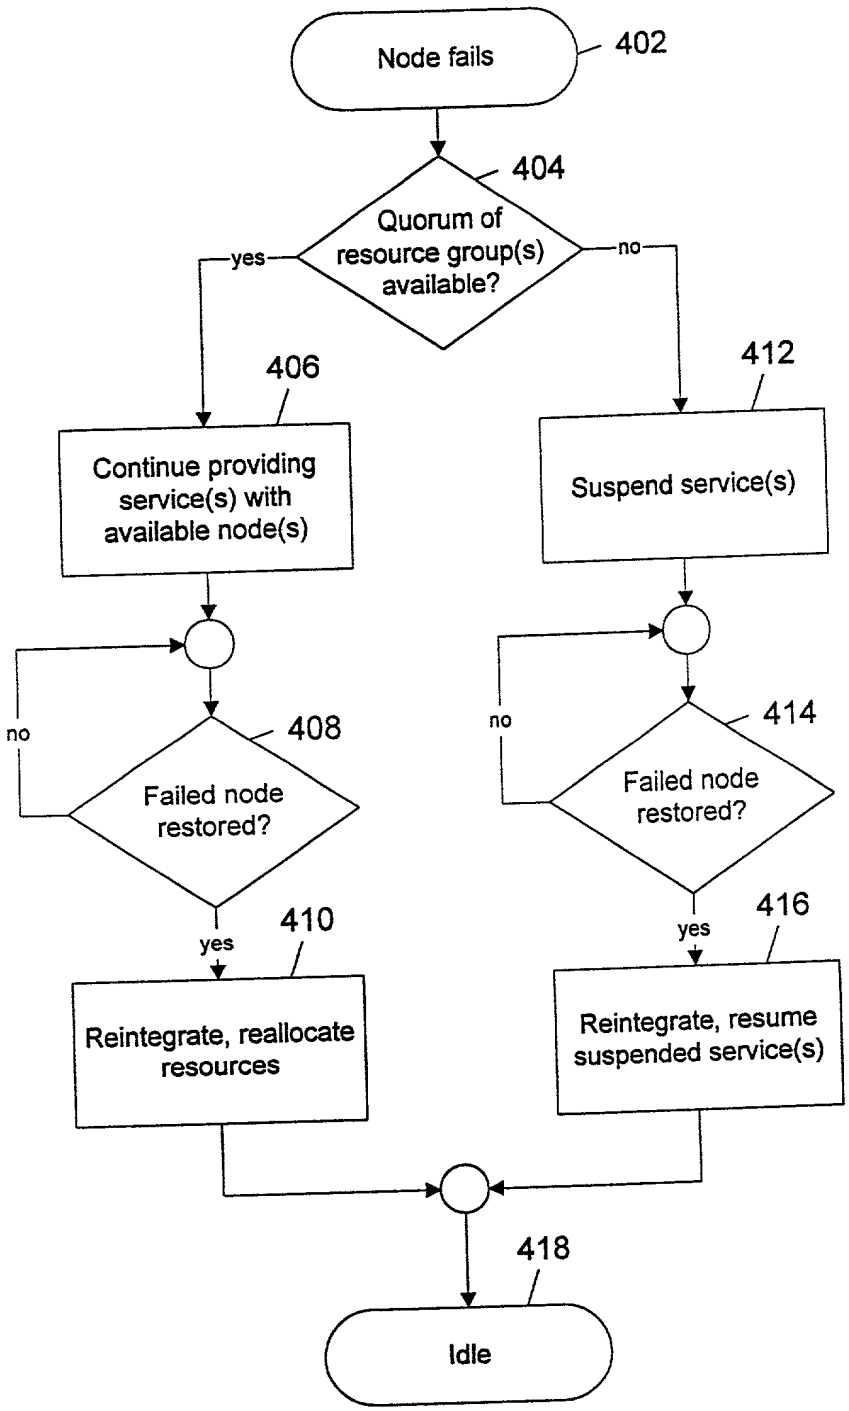
Node fails
402
404
Quorum of resource group(s) available?
yes
no
406
Continue providing service(s) with available node(s)
412
Suspend service(s)
no
408
Failed node restored?
yes
410
Reintegrate, reallocate resources
no
414
Failed node restored?
yes
416
Reintegrate, resume suspended service(s)
418
Idle

Figure 4